arXiv:1609.08372v2 [cs.OS] 28 Sep 2016

An Evaluation of Coarse-Grained Locking for Multicore Microkernels

Kevin Elphinstone, Amirreza Zarrabi, Adrian Danis, Yanyan Shen, Gernot Heiser
UNSW and Data61, CSIRO, Australia
firstname.lastname@data61.csiro.au

Abstract

The trade-off between coarse- and fine-grained locking is a well understood issue in operating systems. Coarse-grained locking provides lower overhead under low contention, fine-grained locking provides higher scalability under contention, though at the expense of implementation complexity and reduced best-case performance.

We revisit this trade-off in the context of microkernels and tightly-coupled cores with shared caches and low inter-core migration latencies. We evaluate performance on two architectures: x86 and ARM MPCore, in the former case also utilising transactional memory (Intel TSX). Our thesis is that on such hardware, a well-designed microkernel, with short system calls, can take advantage of coarse-grained locking on modern hardware, avoid the run-time and complexity cost of multiple locks, enable formal verification, and still achieve scalability comparable to fine-grained locking.

1. Introduction

Waste of processing power resulting from lock contention has been an issue since the advent of multiprocessor computers, and has become a mainstream computing challenge since multicores became commonplace. Much research is directed to understanding and achieving scalability to large numbers of processor cores, where lock contention is inevitable and must be minimised [Clements et al., 2013]. *It is now taken as given that locks must be fine-grained*, ideally protecting individual accesses to shared data structures, and that shared data structures must be minimised, or, in the extreme case of a multikernel [Baumann et al., 2009], avoided altogether.

We observe that a discussion of scalability cannot be done without taking into account operating system (OS) structure as well as platform architecture. Prior scalability work is typically performed in the context of a monolithic OS that needs to scale to hundreds or thousands of concurrent hardware execution contexts, and communication between contexts measuring in the thousands of cycles. But monolithic systems are no longer all that matters, microkernels are finding renewed interest due to their ability to reduce a system's trusted computing base and thus its attack surface [Heiser and Leslie, 2010; Klein et al., 2014; McCune et al., 2008; Steinberg and Kauer, 2010; Zhang et al., 2011].

In a monolithic system, such as Linux, typical system call latencies are long, even compared to inter-core communication latencies in the 1000s of cycles. In contrast, a well-designed microkernel is essentially a context-switching engine, with typical syscall latencies in the hundreds of cycles [Heiser and Elphinstone, 2016]. In such a system, the cost of cross-core synchronisation may be an order of magnitude higher than the basic syscall cost. It therefore makes no sense to run a single kernel image, with shared data structures, across such a manycore machine. An appropriate design should share no data between cores where communication is expensive, resulting in a multikernel design [Baumann et al., 2009].

However, the multikernel approach is not the complete answer either. It presents itself to user-level as a distributed system, where userland must explicitly communicate between nodes. This is not the right model where communication latencies are small, eg. across hardware contexts of a single core, or between cores that share an L2 cache, where they are of the order of tens of cycles, well below the latency of a syscall even in a microkernel. In this context, explicit communication between nodes is more expensive than relying on shared memory, and there is no justification for forcing a distributed-system model on userland.

We therefore argue that, for a microkernel, the right model is one that reflects the structure of the underlying hardware: a shared kernel within a closely-coupled cluster of execution contexts, but shared nothing between such clusters. The resulting model is that of a *clustered multikernel* [von Tessin, 2012].

A node in such a cluster puts scalability into a different context: rather than to hundreds or thousands of cores, it only needs to scale to the size of a closely-coupled cluster, no more than a few dozens of execution contexts. Such a cluster matches another important category of platforms: the now ubiquitous (and inexpensive) low-end multicore processors deployed by the billions in mobile devices.

A microkernel for a closely-coupled cluster represents an area in the design space markedly different from that of manycores, and it is far from obvious that the same solutions apply. In particular, it is far from obvious that fine-grained locking makes sense. In fact, the typical systemcall latencies are not much longer than critical section in other OSes. In that sense, for a microkernel, even a big lock is not much coarser than a fine-grained lock in a system like Linux – as has been observed before, for a microkernel, a big lock may be fine-grained enough [Peters et al., 2015].

This discussion may seem academic at first, given that fine-grained locking techniques are well-known and widely implemented, so why not use them anyway? There are, in fact, strong reasons to stick with coarse-grained locking as long

as possible: Each lock acquisition has a cost, which is pure overhead in the absence of contention. While insignificant compared to the overall system-call cost in a system like Linux, in a microkernel this overhead is significant.

More importantly, the concurrency introduced by fine-grained locking greatly increases the conceptual complexity of code, and thus increases the likelihood of subtle bugs that are hard to find [Lehey, 2001], as painfully confirmed by our experience implementing fine-grained locking in seL4. Furthermore, this complexity is presently a show-stopper for formal verification, which otherwise is feasible for a microkernel [Klein et al., 2014].

Additionally, as Intel TSX *restricted transaction memory* (RTM) extensions become widely available, there is an opportunity to have the complexity of coarse-grained locking and the performance of fine-grained locking by using RTM to elide coarse-grained locks [Rajwar and Goodman, 2001].

We therefore argue that it is important to understand the performance impact of a big-lock design, which maximises best-case performance, minimises complexity and eases assurance. To this end we conduct a detailed examination of the scalability of the seL4 microkernel on closely-coupled clusters on two vastly different hardware platforms (an x86-based server and an ARM-based embedded processor) under different locking regimes. We make the following contributions:

- We estimate the theoretical performance of a coarsely synchronised (big-lock) microkernel using queueing theory (Section 4).
- We validate the queueing model experimentally, and at the same time identify modifications to the microkernel to achieve near theoretically optimum performance (Section 5.1).
- We compare big-lock and fine-grained-lock implementations of the seL4 microkernel and evaluate those on closely-coupled cores on two architectures (ARM and x86), in contrast to the usual approach of aiming for high-end scalability across loosely-coupled cores (Section 5).
- We present (on x86) the first use of hardware transactional memory that places the majority of the kernel into a single transaction for concurrency control, and we compare it with locking (Section 5).

We show that the choice of concurrency control in the kernel is clearly distinguishable for extreme synthetic benchmarks Section 5.1). For a realistic, system-call intensive workload, performance differences are visible, and coarse grain locking is preferred (or the equivalent elided with hardware transactional memory) over extra kernel complexity of fine-grained locking (Section 5.2).

2. Background

2.1. Locking granularity

The best locking granularity is determined by a trade-off involving multiple factors. As long as there is no contention, taking and releasing locks is pure overhead, which is minimised by having just a single lock, the *big kernel lock* (BKL). Each lock adds some overhead which degrades the best-case (i.e. uncontended) performance.

As long as the total number of locks is small, this baseline overhead is usually small compared to the basic system-call cost. However, on a well-designed microkernel, where system calls tend to be very short (100s of cycles) this overhead might matter.

Fine-grained locking can significantly reduce contention, if it enables unlocked execution of the majority of code. In a BKL kernel, contention can be expected to be noticeable as soon as the *hold time* (fraction of time spent inside the kernel, also referred to as *kernel time*) is not small compared to the *pause time* (fraction of time spent in user mode).

The amount of kernel time depends on the profile of system calls executed, and thus on the workload. On a monolithic kernel, most system services are provided by the kernel, especially I/O, and consequently I/O-intensive workloads tend to have high kernel time. On a microkernel, system services are provided by server processes running in user mode, and the kernel provides communication between clients and servers. On a well-designed microkernel, such as the ones of the L4 family, kernel time is dominated by context switches [Liedtke, 1995]. The total number of kernel calls is higher than in a monolithic kernel (at least twice as high, as every server invocation invokes the kernel twice) but the average system-call latency is a tiny fraction of that of a monolithic kernel.

Hence, a BKL is a more credible design for a microkernel than for a monolithic OS, at least for a closely-coupled cluster of execution contexts, where intra-cluster communication latencies are low (eg. due to a shared cache). As explained in the introduction, this does not prevent the kernel from use in manycores, but on such hardware, kernel state should not be shared across clusters, resulting in the clustered multikernel design. It is the design of a kernel for such a cluster which we explore in this paper. To avoid drawing invalid conclusions from idiosyncrasies of a particular platform, we examine two very different architectures: an x86-based server processor and an ARM-based processor aimed at embedded devices, especially smartphones.

2.2. x86 platform

As an x86 platform we use a server-class Dell Poweredge R630 fitted with two Intel Xeon E5-2683 v3 processors. These are a 14-core processors with a base clock rate of 2.0 GHz and two hardware threads each, giving 28 hardware threads per processor. Thus the machine has total of 56 hardware threads across the two CPU sockets. While not officially supported, the microarchitecture features Intel's TSX implementation of *restricted transactional memory* (RTM), which we describe further in Section 2.4.

The processor features three levels of cache. Each core has private L1 instruction and data caches, each 32 KiB in size

Platform	Memory Level	Local core	Intra-socket	Inter-socket
x86	L1	4	115	218
	L2	12	105	208
	L3	44	44	163
	Memory	185	185	265
ARM	L1	4	17	N/A
	L2	26	28	N/A
	Memory	140	N/A	N/A

Table 1: Memory and cache access latency in cycles.

and 8-way associative. Each core furthermore has a private, non-inclusive, 8-way 256 KiB L2 cache. The last-level cache is 35MiB, consisting of a 2.5MiB slice per core.

Table 1 shows our measured memory latency, cache access latency, and latency of data transfer between cores on the same socket, and across separate sockets. The measurements were obtained using code derived from *BenchIT*, and the results are reasonably consistent with measurements obtained by the authors [Molka et al., 2015], noting the differing clock rates of the system under test. One should also note that these results are sensitive to the distance between cores and thus will vary depending on the specific cores involved.

2.3. ARM platform

Our ARM platform is the Sabre Lite, which is based on a Freescale i.MX 6Q SoC, featuring a quad-core ARM Cortex-A9 MPCore processor [Freescale, 2013].

The cores run at a 1 GHz clock rate and have private, split L1 caches, each 4-way-associative and 32 KiB in size. The cores share a 1 MiB, unified, 16-way-associative L2 cache, which is the last-level cache. We ported the microbenchmarks from *BenchIT* to the ARM platform and obtain the results in Table 1.

Compared to x86, the ARM has much lower latency of data transfer between caches, and the latency is unaffected by distance between the cores.

2.4. Intel TSX

TSX provides 4 new instructions: `XBEGIN`, `XEND`, `XTEST` and `XABORT`. Code successfully executed between `XBEGIN` and `XEND` instructions will appear to have completed atomically, and is thus called a transactional region. If there are any memory conflicts during the execution of the transactional region, the transaction will abort and jump to the instruction specified by the `XBEGIN`. A program can explicitly abort a transaction by issuing an `XABORT` instruction. `XTEST` returns whether currently executing within a transactional region.

TSX takes advantage of existing cache coherency protocols, to identify sets of cache lines written to and read by different cores on the CPU. This has two important consequences: memory conflicts are captured at a cache-line granularity, and transactions are constrained by the size of the L1 and L2 caches. The mutated state must fit inside L1 cache, and the accessed state must fit inside the L2 cache [Hasenplaugh et al., 2015]. The consequence is that it is probably not feasible to wrap a complete monolithic kernel into an RTM transaction, as it is unlikely to fit within the L1 and L2 caches.

Owing to the implementation of TSX, the RTM lock logically protects a dynamic set of individual L1 and L2 cache lines, and as such is a fairly extreme case of fine-grained locking, which should result in much reduced contention (assuming a sane layout of kernel data structures).

Note that an RTM transaction is not guaranteed to complete, even when the transaction is small enough and has no memory conflicts. A variety of (hardware-implementation specific and frequently unspecified) scenarios can result in an abort. Of particular interest to our work are certain interactions on specific registers that trigger aborts, but are clearly unavoidable when executing OS code.

Given transactions have no guarantees of progress, the developer must ensure that there exists a fallback method of synchronisation that ensures progress in the presence of repeated aborts. We use the commonly implemented technique of falling back on a regular lock for the code fragment in the case of repeated aborts. To avoid races between transactions and locks, our transactions test the lock upon entry to an RTM section, to ensure the lock is free and force it into the read set of the transaction. A change in lock state by a competing thread will trigger the desired abort, and allow the section to synchronise via the lock.

3. Microkernel Implementation

3.1. seL4

As we use seL4 as our microkernel testbed, we will now summarise its relevant features, Klein et al. [2014] presents more details. seL4 is event-based, with a single kernel stack. To aid verification, seL4 uses a two-phase system call structure, where the first phase confirms the pre-conditions required for system call execution, and the second phase executes the system call without failure. Blocking operations are handled by re-starting the system call and thus re-confirming the preconditions prior to continuing execution.

The kernel executes with interrupts disabled. This concurrency-free design has traditionally been used in L4 kernels in order to achieve high best-case performance, and has been used on other systems as well [Ford et al., 1999]. With formal verification it becomes as necessity, as it is for now infeasible to verify concurrent kernel code.

The kernel features some long-running operations resulting from the destruction of kernel objects that may have derived objects. In order to achieve usable interrupt latencies, it has explicit preemption points, where the kernel polls for pending interrupts, and restarts the operation if there are any [Blackham et al., 2012]. The restart allows interrupts to be triggered from outside the kernel, prior to continuing the original operation.

seL4 supports the traditional L4-style synchronous (rendezvous) message passing IPC with a payload of up to a few hundred bytes. IPC operates via port-like objects called *endpoints*. In addition, the kernel provides *notifications* with semantics similar to binary semaphores.

3.2. Big kernel lock

The BKL is the natural, minimal extension of the existing seL4 design to multicores, as it is easy to implement and mostly preserves the in-kernel assumption of no concurrency. The kernel entry and exit code, which saves and restores the user-state to a per-core kernel stack and sets up safe kernel execution, remains outside of the BKL, while the rest of the kernel is protected by the BKL.

This design is not entirely sufficient – the following invariant, used in the verification, no longer holds on a multicore kernel, even when the BKL is held:

> Except for the currently executing thread's TCB and page table, all other TCBs and page tables are quiescent, and can be mutated or deleted.

User-level code executing on other cores implicitly depends on the running thread's TCB and page table to transition to kernel-mode via the kernel entry code to compete for the BKL. The invariant therefore no longer holds. We address this by modifying the kernel to ensure remote cores are not dependent on any TCB or page-table undergoing deletion.

We modify our prototype to keep a bitmap of cores that have seen a specific page table in the page table itself, and IPI only those cores to trigger the remote core to enter the kernel idle loop (which has a permanently allocated TCB and page-table), and also to shoot down the TLB. A TCB can be identified as active via the CPU affinity in the TCB itself combined with the per-core *current thread* pointer of the remote core, in which case the TCB is handled in a similar manner to the page table.

This design, which is partially driven by the existing event-driven code base, is a valid design choice thanks to the short duration of most system calls in the microkernel; it would result in poor scalability on any other kind of system.

The only other required change is introducing per-core idle threads. However, in order to minimise inter-core cache-line migrations, we also introduce per-core scheduler queues in addition to the current-thread pointers, even though access is serialised by the BKL. This partitioned scheduling implies that threads can only migrate between cores if explicitly requested by the user, which is consistent with seL4's general philosophy of having all resource management under user control (and also helps reasoning about real-time properties).

To reduce contention (and enable the use of transactional memory, see Section 2.4) we further minimise the amount of locked code by moving context-switch-related hardware operations after the BKL release, which has the benefit of reducing the critical section length.

We use a CLH lock, as scalable queue lock [Craig, 1993], to synchronise the BKL kernel variant.

3.3. Fine-grained locking

To compare the coarse-grained BKL with more complex but more scalable fine-grained locking, we first replace the BKL with a big reader lock [Corbet]. The lock allows all reader cores to proceed in parallel as they access only local state to obtain a read lock.

In our present prototype we use a single write lock around the the non-IPC-related kernel code paths. These code paths, generally dealing with resource management, are infrequently executed, compared to IPC and interrupt handling, and as such not performance critical. This allows us to avoid significant code changes without affecting overall performance.

This design allows us to gradually migrate the kernel code out of the writer lock into the reader lock. As long as deallocation of kernel objects remains inside the writer lock, memory safety is retained while holding the reader lock. Moving code into the reader lock exposes the *contents* of the objects to concurrency for improved scalability, which can then be protected using individual fine-grained locks.

IPC mutates the state of TCBs, endpoints, and (potentially) the scheduler queues (depending on whether optimisations apply that avoid queue updates during IPC [Heiser and Elphinstone, 2016]). We add ticket locks to each of these data structures for synchronising IPC within the reader lock. A typical IPC now involves the kernel reader lock, two TCB locks, and one endpoint lock. Lock contention during IPC is limited to cases where IPC involves a shared destination or endpoint, or general contention with the kernel writer lock. Independent activities performing IPC on independent cores result in no lock contention. We avoid deadlocks by identifying the affected TCBs prior to locking (made possible by memory safety provided by the reader lock), and then locking them in order of their memory addresses.

3.4. Hardware transactional memory

The TSX extensions, combined with the small size of the kernel, allow us to optimistically execute the majority of the code without concurrency control. This is analogous to taking the BKL kernel variant and speculatively eliding the lock [Rajwar and Goodman, 2001]. The event-based design of the kernel is an important enabler for lock elision as it avoids blocking. We bracket almost the entire kernel with the transaction primitives shown in Figure 1. The somewhat simplified code is self explanatory, except the 'L' argument to `_xabort()`, which is returned as the status at `_xbegin()` to distinguish between abort types.

In addition to the changes described in Section 3.2, we need to move any TSX-specific abort-triggering CPU operations after the transaction. Many of those do not occur in seL4, as most aborting operations are typical for device drivers, which

```
beginTransaction() {
  while ((status = _xbegin()) !=
         _XBEGIN_STARTED ) {
    txnAttempts++;
    if (txnAttempts >=
        RTM_ATTEMPTS_THRESHOLD) {
      break; /* Give up */
    }
    /* wait for lock freed before retrying txn */
    while(LockTest());
  }
  if (status == _XBEGIN_STARTED) {
    if (LockTest()) {  /* not free */
      _xabort('L');
    }
  } else {
    lockAcquire(); /* BKL fall-back */
  }
}

endTransaction() {
    if ( (txnInside = _xtest()) ) {
        _xend();
    } else {
        lockRelease();
    }
}
```

Figure 1: Kernel transaction pseudo code.

are user-level programs in seL4. The remaining problematic operations are:

- context-switch-triggered page-table register (CR3) loading and segment-register loading;
- IPI triggering for inter-core notifications;
- interrupt management for user-level device drivers, which consists of masking and acknowledging interrupts prior to return to the user-level handler.

The key insight here is that it is safe to move these operation outside of the transaction, because the two-phase kernel ensures the system call which requires these operations is guaranteed to succeed once the execution phase is entered, and that these operations are local to a core and thus are not exposed to concurrent access from other cores. Note that preemptions during this code section are prevented since the kernel runs with interrupts disabled.

4. BKL multicore scalability

In this section we use queuing theory to model the scalability of a BKL microkernel. The theoretical model provides us with a method to estimate best-case performance for a workload parametrised by the rate the lock can be serviced (μ), the arrival rate of the lock requests (λ, i.e. system call rate), and the number (n) of cores in the machine. An estimate of best-case performance provides a theoretical reference point to target.

4.1. Modelling contention

We employ the *machine repairman* queuing theory model. The model is historically based on machine failures in a factory (characterised by a failure rate) combined with waiting for a single repairman (the service rate). In our case the model corresponds to the arrival rate of lock requests combined with the service rate of the lock itself.

The model for an n-core multiprocessor has $n+1$ states, representing the number of cores holding or waiting to acquire the kernel lock, as illustrated in Figure 2.

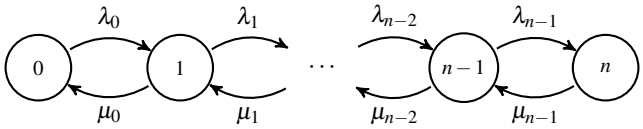

Figure 2: The machine repairman model.

The model assumes the rate, μ, of servicing the lock, to be independent of the number of cores queued, i.e. $\mu_k = \frac{1}{s}$, where s is the average service time of the lock. It further assumes that the rate of arrivals, λ, is proportional to the number of cores not already waiting for the lock, i.e. $\lambda_k = \frac{n-k}{a}$, where a is the average inter-arrival time for a single core in the absence of contention.

In a steady state, the rates of lock acquisitions and lock releases must be balanced. If P_i is the probability of being in state i, this means $P_k \cdot \lambda_k = P_{k+1} \cdot \mu_k$. From this we can derive the probabilities as

$$P_k = P_0 \cdot \frac{s^k n!}{a^k (n-k)!} \,. \tag{1}$$

The system must always be in one of these states, $\sum_{i=0}^{n} P_i = 1$, from which we can obtain

$$P_0 = \frac{1}{\sum_{i=0}^{n} \left(\frac{s^i n!}{a^i (n-i)!} \right)} \,, \tag{2}$$

and ultimately

$$P_k = \frac{\frac{s^k}{a^k (n-k)!}}{\sum_{i=0}^{n} \left(\frac{s^i}{a^i (n-i)!} \right)} \,. \tag{3}$$

From this we can compute the expected queue length as

$$w = \sum_{i=0}^{n} i P_i \,, \tag{4}$$

and lock throughput is

$$\mu (1 - P_0) = \frac{1 - P_0}{s} \,. \tag{5}$$

4.2. Model Assumptions and Kernel Design

The queueing model assumes that the average rate of serving the lock is independent of the queue depth. This is not true for non-scalable locks or in the case of mutating shared kernel state [Boyd-Wickizer et al., 2012]. We satisfy these assumptions by avoiding shared mutable state for unrelated kernel system calls through per-core data structures, and using the scalable CLH lock.

In addition, peak throughput is inversely proportional to lock service time. Hence, moving as much code out of the lock as possible, in particular the expensive local hardware operations (such as triggering of IPIs, and page table register updates), will improve scalability.

5. Evaluation

To evaluate our multicore microkernel variants, we use two IPC microbenchmarks and a server-style macrobenchmark, as described in the following sections. The platforms under test have been already been described in Sections 2.2 and 2.3.

5.1. Microbenchmarks

5.1.1. Single-Core IPC microbenchmarks IPC performance is a key contributor to overall system performance in microkernel-based systems, and optimising IPC performance has a long history in the L4 community [Heiser and Elphinstone, 2016]. The traditional benchmark for best-case IPC performance is "ping-pong": a pair of threads on a single core does nothing other than sending messages to each other. This allows us to asses the basic cost of our lock implementations, i.e. the pure acquisition and release cost, without any contention.

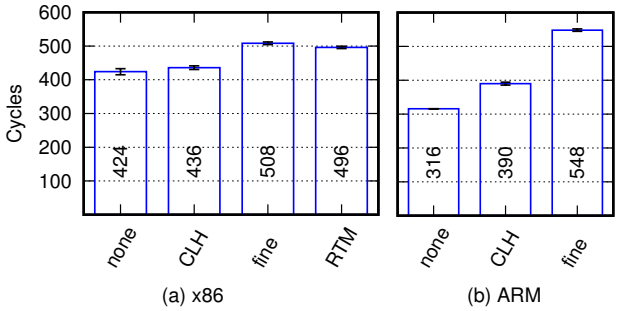

Figure 3: Raw one-way IPC cycle cost for different seL4 locking mechanisms. Error bars indicate standard deviations.

The figure shows that on x86, the overhead of a single CLH lock (BKL) is approximately 3% compared to a baseline uniprocessor kernel with no concurrency primitives ("none"). With fine-grained locking, however, the overhead is 20%. The overhead of uncontended transactions is 17%.

On the ARM, the cost of a single lock is significantly higher with 23%, while for fine-grained locking it is over 70%. The higher synchronisation costs on the ARM processor relate to its partial-store-order memory model. It requires memory barriers (`dmb` instructions) to preserve memory-access ordering. In our experience, the barriers cost from 6 cycles up to 19 cycles depending on micro-architectural state. Our implementation of CLH executes 6 barriers on this benchmark, while 16 are needed with fine-grained locking. These barriers explain most of the overhead.

The significant cost of fine-grained locking provides a motivation for sticking with the BKL as long as possible, even if verification tractability was no issue.

5.1.2. Multicore IPC BKL microbenchmarks To explore scalability and experimentally validate the queueing model, we extend the single-core ping-pong to multiple cores. Specifically, we run a copy of ping-pong on each hardware thread, with all hardware threads executing completely independently and unsynchronised. We use the BKL kernel on the x86 platform.

We add an exponentially distributed random delay between receiving and replying to IPC, for each ping-pong pair. The delay varies between an average of 500 and 32,000 cycles, in powers of 2, to create seven individual microbenchmarks that simulate a work load from an extreme system call intensive workload to a relatively compute-bound workload.

The benchmark is *embarrassingly parallel* ensuring that limits of scalability are related to our kernel design and implementation, and not the benchmark itself. The number of hardware threads (cores in this case) varies between 1 and 28.

This benchmark produces extreme contention on the kernel (for low delay values). However, *none of the kernel data structures are contended*, as each hardware thread's pair of software threads accesses disjoint kernel objects (TCBs and IPC endpoints) during their syscalls. Hence, while expecting contention on the BKL, fine-grained locking and RTM can be expected to scale perfectly.

For each delay and core-count parameter pair, the benchmark consists of a two second warm-up, followed by sampling total IPCs during a one second interval to give total IPC throughput per second.

Figure 4(a) plots the resulting overall IPC throughput for varying number of hardware threads. Each point represents one measurement for a particular delay time (identified by symbol and colour). The vertical dotted lines shows where the cores are split across the two CPU sockets. The results have negligible variance.

For the runs with an average delay of 2000 cycles we perform a least-squares regression of the queueing model using only the points for the first 14 hardware threads (i.e. within one socket). The regression yields a service time of 358 cycles, and an average delay of 1999 cycles, with $R^2 = 0.99$, meaning that the intra-core results are explained by the model if the service time is 358 cycles.

We use this service time to predict throughput for all other values of the delay parameter, resulting in the solid lines in

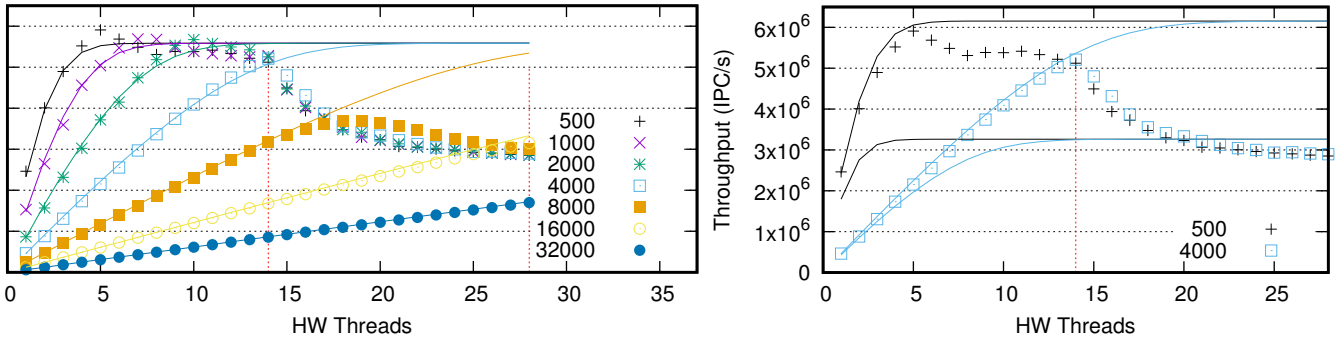

(a) Scalability measurements and fits for 358 cycles service time.

(b) Scalability bound prediction for service times of 323 cycles (upper lines) and 613 cycles (lower lines).

Figure 4: Total IPC throughput for varying parallelism and delay times (cycles). Points are measurements, lines are queueing model fits.

the graph. We can see that in all cases, these fit the observed throughput values very well for at least 14 threads (i.e. the model explains the intra-socket behaviour well).

The model breaks down once cores of the second socket are involved, except for the highest delay times. This is unsurprising, as with multiple sockets, the assumption of a fixed service time no longer holds, as transfer times for the cache line holding the lock now depend strongly on locality. We confirm this by instrumenting the lock: The average holding time is 164 cycles for a single core. For four or more cores, the observed holding times vary between 323 and 613 cycles.[1]

In Figure 4(b) we repeat the experimental results of Figure 4(a) for two delay values, 500 and 4,000 cycles. We also show the model prediction for 323 and 613 cycles, the minimum and maximum holding times we observed for thread counts of four or more. We can see that the results are upper-bounded by the 323-cycle curve, and, as long as only one socket is used, remain close to the bound. Once the second socket is used, the lines quickly approach the lower line, corresponding to the higher service time, which remains a lower bound.

5.1.3. Observations about the model The queueing model accurately predicts experimental results where the average lock holding time is stable. Where the lock holding time is variable, it can be used to predict a performance range. The model enables prediction of where the *knee* of the performance curve occurs for a given lock holding time and average delay between system calls, assuming the absence of other application-related limiting factors.

The lock-holding time range for the microkernel varies from an average 164 cycles on a single core, to approximately 300 cycles within a socket, to 600 cycles distributed across sockets. Thus the lock holding time is dominated by the architectural cost of cache line transfer for the lock. This is indirect confirmation that our microkernel is indeed scalable in the sense of not sharing any mutable state across cores except for the lock itself. It also implies that any improvement in reducing lock holding time on a single core will have only a modest effect on overall scalability due to the high architectural costs on the Xeon.

The model and experiments show that a workload running on the microkernel with an average delay between system calls on each core of 4000 cycles would scale to 14 cores, i.e. a single CPU socket. An average delay time of 16000 cycles is needed for a workload to scale across both sockets. The results support our hypothesis that a big lock will scale as long cores are closely coupled.

These results are readily applicable in general to conservative locks protecting *potentially* contended data, that rarely contend in practice. We also note for the following section that these results are readily applicable to the abort path in the RTM variant of the microkernel.

5.1.4. Locking variant evaluation To compare different lock variants, we run on x86 with a 4,000-cycle average delay between system calls, as this is just above the scalability limit of the BKL across a whole socket.

Figure 5(a) shows that RTM behaves identically to fine-grained locking. This is expected: as explained in Section 2.4, RTM is logically an extreme case of fine-grained locking, and the baseline lock overhead is the same as for the fine-grained locks according to Figure 3.

The BKL variant serialises the IPC path across all the available hardware threads, thus hits the knee in the performance curve as predicted by the queueing model. As also expected, it converges on a lower performance plateau as the benchmark spans the sockets.

On ARM (Figure 5(b)), where intra-core cache line migration costs are very low (in terms of cycles), we chose an agressive 500-cycle average delay. In addition, we run the pathological zero-delay case, where the system call rate is only limited by the user-level stubs and cost of the system call

[1] The average holding time of the lock is strongly influenced by the average transfer time of the cache line. Given transfer cost is zero for re-acquisition of the lock by the same core, low core counts have higher probability of re-acquisition, and thus unrealistically low average holding times for the general case.

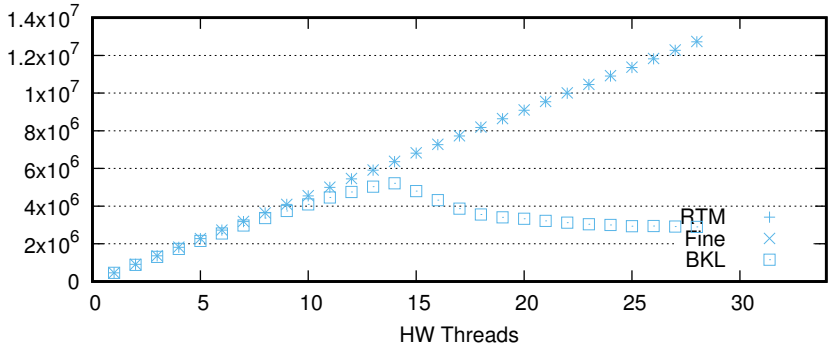

(a) x86, average delay 4,000 cycles.

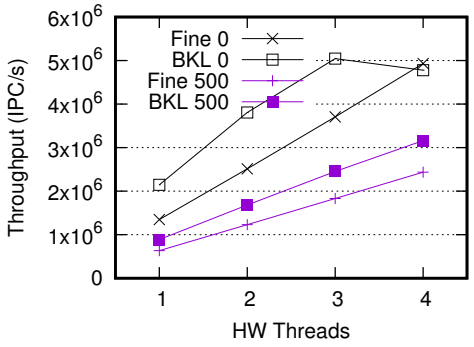

(b) ARM, average delay 0 and 500 cycles. Lines are for clarity (no fit).

Figure 5: Total IPC throughput for varying parallelism and different locking implementations.

itself. We see the higher overhead for fine-grained locking is readily visible, with the BKL variant outperforming the fine-grained variant for the fairly extreme case of the 500 cycle average delay time, with perfect scalability to all four cores. It takes the unrealistic minimal delay for the BKL variant to plateau at 3 cores, allowing the fine-grained variant to exceed BKL throughput. We can expect the BKL to scale significantly beyond the size of our quad-core machine for realistic workloads.

5.2. Redis-based Macrobenchmark

In order to assess BKL scalability, and the significance of the overheads of the fine-grained schemes, we look for a "realistic worst-case" scenario, i.e. a benchmark which produces as high as system-call rate as can be expected under realistic conditions.

None of the usual embedded-system benchmarks produce significant syscall loads on the microkernel, we therefore use a server-style benchmark. Note that the nature of the benchmark is completely irrelevant for this exercise, all that counts is the rate and distribution of kernel entries. The relevant operations are IPC and interrupt handling, as all other microkernel operations deal with resource management that is relatively infrequent.

The seL4 equivalent of a syscall in a monolithic system is sending an IPC message to a server process and waiting for a reply (i.e. two microkernel IPCs per monolithic OS syscall). Similarly, an interrupt, which in a monolithic OS results in a single kernel entry, produces two for the microkernel-based system, as the interrupt is converted by the kernel into a notification to the driver (one kernel entry), and the driver acknowledges to the kernel with another syscall.

5.2.1. Benchmark setup In order to hammer our kernel, we use a simple client-server scenario, consisting of the Redis key-value store [Redis]. We consolidate the clients and servers on the same machine due to insufficient network bandwidth to saturate the large number of cores. Redis receives client requests from a virtual network processor on Core 0. Each client and server has their own private copy of the lwIP TCP/IP stack [lwIp] running as a usermode process.

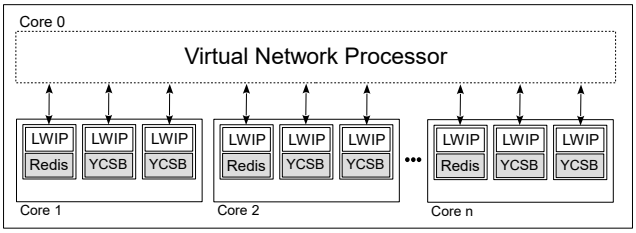

Figure 6: Redis-based benchmark architecture.

Figure 6 shows the system under test. With the exception of Core 0, each core has a Redis server and two copies of the Redis benchmarking client.

We run Redis as volatile instances, i.e. disabling file-system access, in order to maximise throughput and therefore the rate of kernel entries.

We evaluate the performance using a modified version of Yahoo! Cloud Serving Benchmarks (YCSB) [Cooper et al., 2010] as the benchmarking client. All client instances start simultaneously and are tuned to perform a fixed number of operations that result in at least 2 minutes of run time. The benchmarking client consists of the read-only workload with zipfian distribution as presented in Cooper et al. [2010]. For each kernel variant, we instrument the kernel to record idle time within the idle loop for obtaining CPU utilisation for each run.

5.2.2. Results Figure 7 shows the results of the Redis benchmark on the x86 platform.[2] Not surprising, given the extreme workload, scalability is limited.

The transactions-based kernel performs consistently best. However, the BKL keeps up until 5 cores, after which throughput starts to drop. Fine-grained locking consistently performs at about 60-75% of the transactions kernel.

[2]These are single runs, sorry, no standard deviations...

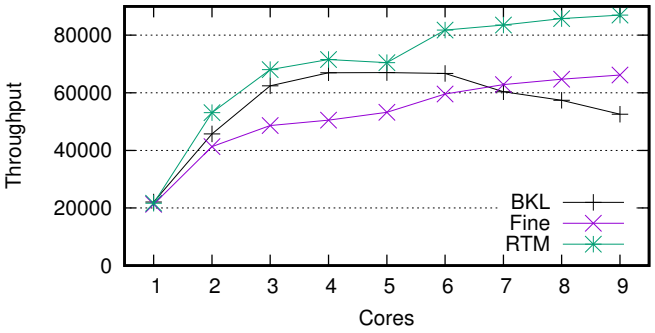

Figure 7: Redis-based benchmark architecture.

We measure the average delay time on each core by dividing user time by the number of kernel entries. We find that the average time is 800 cycles on core 0 (the virtual network), and around 1,600 cycles on the client/server cores (figures taken from the five-core case). This confirms that the benchmark is fairly extreme in terms of system call load.

We can relate the observed delay times back to Figure 4(a). With 800 cycles delay on one core and 1,600 on the others, we expect a behaviour similar to the 1,000-cycle delay curve of Figure 4(a). And the similarity is indeed there: the 1,000-cycles curve peaks at 7 cores and then drops of slowly. The BKL curve of Figure 7 peaks slightly earlier but overall looks similar.

We re-iterate that this is an extreme benchmark, and Figure 4(a) tells us that a slightly less extreme version, with about four times the average delay, should scale to a full socket.

6. Related Work

Writing parallel and scalable code is a topic almost as old as computing itself. Cantrill and Bonwick [2008] provide some historical context and motivation for concurrent software, together with words of wisdom to tackle the difficulties of writing high-performance and correct concurrent software. We adhere to their advice by avoiding parallelising complex software (i.e. splitting the BKL) as our data shows it is unwarranted for closely coupled cores.

Recent complementary work evaluates the scalability of various synchronisation primitives [David et al., 2013] on many-core processors. The authors reinforce that scalability is a function of the hardware, with scalability best when access is restricted to a single socket with uniform memory access – exactly our area of interest.

Boyd-Wickizer et al. [2012] use queueing theory to predict the collapse of ticket locks in Linux. We use a different model to predict performance of a microkernel synchronised with a scalable lock, not just the lock itself. We also use it to validate the scalability of the implementation of the microkernel itself.

Hardware transactional memory is utilised in TxLinux [Rossbach et al., 2007] to implement *cxspinlocks*, a combination of co-operative spinlocks and transactions capable of supporting device I/O and nesting. A small microkernel needs neither, as I/O is at user-level, and it can be designed to avoid complex, nested, fine-grained locks.

Hofmann et al. [2009] apply HTM to coarse-grained locks in Linux 2.4 on a simulated HTM system. Our goals are similar, however our experiments are on real hardware (Intel TSX), on a microkernel with a single lock. Our event-based kernel avoids the need for a *cxmutex* to handle waking waiting threads on exiting a transaction.

Patches for Linux to utilise Intel TSX have been made available [Kleen]. To our knowledge, no performance data was released. Eliding existing fine-grained locking does nothing to reduce kernel complexity. We elide the whole microkernel, providing favourable performance while retaining simplicity.

7. Conclusions

We have analysed scalability of a microkernel with fast system calls across closely-coupled processor cores. We find that for such a system, the overhead of locking is significant, ranging between a few percent for a single lock on x86 to 23% on an ARM processor. This makes the best approach to concurrency control non-obvious, especially when keeping in mind that it makes no sense to scale such a kernel to a large multicore, where inter-core cache-line migration latencies exceed the basic syscall cost.

We analysed three different multicore implementations of seL4, one using a big kernel lock, one using fine-grained locking, and the third using hardware-supported transactions. We evaluated the implementation on a server-class x86 platform as well as an ARM multicore aimed for the embedded market. We support the experiments with a queueing-theoretical model.

There are three main take-aways from our evaluation. One is that the inter-core cache-line migration cost matters a lot. This is demonstrated by the ARM results, where the BKL scales perfectly to 4 cores even with the unrealistically high 500-cycle average inter-syscall time. If architects can maintain similarly tight coupling with higher core counts, our modelling shows that the BKL can be expected scale further.

In contrast, the x86 platform shows that the perfect scaling regime does not extend past two cores even with about double the inter-syscall time, and plateaus at four cores (but keep in mind that this is still an extremely, if not unrealistically high load).

The second take-away is that lock overhead is significant on a well-designed microkernel with very short syscall latencies. This is particularly obvious on the ARM with its relaxed memory model and resulting high barrier costs, but the effect is also significant on the x86, where the kernel using fine-grained locking performs only at about 75% of the BKL version until the latter reaches its performance knee.

The third takeaway is that hardware transactions are an exciting development. To our knowledge, we are the first to implement lock elision for a BKL kernel using Intel's RTM. We show in microbenchmarks that a theoretically embarrassingly parallel application scales perfectly with little overhead

and no serialisation. In our realistic (but extreme) macro benchmark, which is less parallel, RTM upper-bounds the performance of both the BKL as well as fine-grained locking. In the case of the microkernel, where the whole system call can be packed into a single transaction, RTM gives get the best of both worlds.

We can summarise our experience that transactions are the way to go, if they are available. Failing that, the big lock is actually a good choice for a fast microkernel, as it is only outperformed by fine-grained locking under extreme circumstances, at least on the kind of closely-coupled system where a single, shared kernel instance makes sense. Under those circumstances, the reduced (compared to fine-grained locking) implementation complexity is a strong asset, as it enables formal verification, which is presently unfeasible for systems using fine-grained locking. The significantly better performance under less extreme workloads is an added benefit.

References

Andrew Baumann, Paul Barham, Pierre-Evariste Dagand, Tim Harris, Rebecca Isaacs, Simon Peter, Timothy Roscoe, Adrian Schüpbach, and Akhilesh Singhania. The multikernel: A new OS architecture for scalable multicore systems. In *SOSP*, Big Sky, MT, US, Oct 2009.

Bernard Blackham, Yao Shi, and Gernot Heiser. Improving interrupt response time in a verifiable protected microkernel. In *EuroSys*, pages 323–336, Bern, Switzerland, Apr 2012.

Silas Boyd-Wickizer, M. Frans Kaashoek, Robert Morris, and Nickolai Zeldovich. Non-scalable locks are dangerous. In *2012 Ottawa Linux Symp.*, Ottawa, CA, Jul 2012.

Bryan Cantrill and Jeff Bonwick. Real-world concurrency. *ACM Queue*, 6(5), Sep 2008.

Austin T. Clements, M. Frans Kaashoek, Nickolai Zeldovich, Robert T. Morris, and Eddie Kohler. The scalable commutativity rule: Designing scalable software for multicore processors. In *SOSP*, pages 1–17, Farmington, PA, US, Oct 2013.

Brian F. Cooper, Adam Silberstein, Erwin Tam, Raghu Ramakrishnan, and Russell Sears. Benchmarking cloud serving systems with YCSB. In *SoCC*, Indianapolis, IN, US, Jun 2010.

J. Corbet. Big reader locks. http://lwn.net/Articles/378911/.

Travis S. Craig. Building FIFO and priority-queuing spin locks from atomic swap. Technical Report UW-CSE-93-02-02, Department of Computer Science and Engineering, University of Washington, 1993.

Tudor David, Rachid Guerraoui, and Vasileios Trigonakis. Everything you always wanted to know about synchronization but were afraid to ask. In *SOSP*, pages 33–48, Farmington, PA, US, Nov 2013.

Bryan Ford, Mike Hibler, Jay Lepreau, Roland McGrath, and Patrick Tullmann. Interface and execution models in the Fluke kernel. In *3rd OSDI*, pages 101–115, New Orleans, LA, US, Feb 1999. USENIX.

Freescale. *i.MX 6Dual/6Quad Applications Processor Reference Manual*, rev. 1 edition, Apr 2013.

William Hasenplaugh, Andrew Nguyen, and Nir Shavit. Quantifying the capacity limitations of hardware transactional memory. In *7th Workshop on the Theory of Transactional Memory (WTTM 2015)*, Jul 2015.

Gernot Heiser and Kevin Elphinstone. L4 microkernels: The lessons from 20 years of research and deployment. *Trans. Comp. Syst.*, 34(1):1:1–1:29, Apr 2016.

Gernot Heiser and Ben Leslie. The OKL4 microvisor: Convergence point of microkernels and hypervisors. In *APSys*, pages 19–24, New Delhi, India, Aug 2010.

Owen S. Hofmann, Christopher J. Rossbach, and Emmett Witchel. Maximum benefit from a minimal HTM. In *14th ASPLOS*, pages 145–156, 2009.

Andi Kleen. RFC: Kernel lock elision for TSX. https://lkml.org/lkml/2013/3/22/630.

Gerwin Klein, June Andronick, Kevin Elphinstone, Toby Murray, Thomas Sewell, Rafal Kolanski, and Gernot Heiser. Comprehensive formal verification of an OS microkernel. *Trans. Comp. Syst.*, 32(1):2:1–2:70, Feb 2014.

Greg Lehey. Improving the FreeBSD SMP implementation. In *2001 USENIX ATC, FREENIX Track*, Boston, MA, US, Jun 2001.

Jochen Liedtke. On μ-kernel construction. In *SOSP*, pages 237–250, Copper Mountain, CO, USA, Dec 1995.

lwIp. Download site. http://www.nongnu.org/lwip/.

Jonathan M. McCune, Bryan Parno, Adrian Perrig, Michael K. Reiter, and Hiroshi Isozaki. Flicker: An execution infrastructure for TCB minimization. In *3rd EuroSys*, Apr 2008.

D. Molka, D. Hackenberg, R. Schöne, and W. E. Nagel. Cache coherence protocol and memory performance of the Intel Haswell-EP architecture. In *2015 Int. Conf. Parall. Processing*, pages 739–748, Sep 2015.

Sean Peters, Adrian Danis, Kevin Elphinstone, and Gernot Heiser. For a microkernel, a big lock is fine. In *APSys*, Tokyo, JP, Jul 2015.

R. Rajwar and J. R. Goodman. Speculative lock elision: enabling highly concurrent multithreaded execution. In *34th MICRO*, pages 294–305, Dec 2001.

Redis. Download site. http://redis.io.

Christopher J. Rossbach, Owen S. Hofmann, Donald E. Porter, Hany E. Ramadan, Bhandari Aditya, and Emmett Witchel. TxLinux: Using and managing hardware transactional memory in an operating system. In *SOSP*, Stevenson, WA, US, Oct 2007.

Udo Steinberg and Bernhard Kauer. NOVA: A microhypervisor-based secure virtualization architecture. In *5th EuroSys*, pages 209–222, Paris, FR, Apr 2010.

Michael von Tessin. The clustered multikernel: An approach to formal verification of multiprocessor OS kernels. In *2nd Workshop on Systems for Future Multi-core Architectures*, pages 1–6, Bern, Switzerland, Apr 2012.

Fengzhe Zhang, Jin Chen, Haibo Chen, and Binyu Zang. CloudVisor: retrofitting protection of virtual machines in multi-tenant cloud with nested virtualization. In *SOSP*, pages 203–216, Cascais, Portugal, 2011.